Importance of Ultrafast Computing for NASA Missions

(Langley Reconfigurable Hypercomputing Research)

for

Dr. Dennis M. Bushnell, Langley Senior Scientist, OD

by

Dr. Olaf O. Storaasli (ACMB, SMC)

Dr. Robert C. Singleterry (ACMB, SMC)

Dr. Jaroslaw Sobieski (ACMB, SMC)

David K. Rutishauser (CSOB, ASC)

August 6, 2002

Questions Addressed:

1. How does ultrafast computing research help Langley & NASA missions?
2. Why NASA? Won't others do this research (100TFLOPs at DOE)?
3. Can Langley make an impact when industry spends $100B/yr on IT research?
4. How does this relate to nano- and quantum computing?

"***In broad terms, our mandate is to pioneer the future . . . to push the envelope . . . to do what has never been done before***." *NASA Vision by Administrator Sean O'Keefe April 12, 2002, Maxwell School of Syracuse University*

Preface

In the spirit of the NASA Vision as espoused by Administrator O'Keeke, we are definitely "pushing the envelope" in our innovative Hypercomputing research. However, it is important to step back and evaluate the importance of our research to NASA Langley, NASA and our country in relation to other progress taking place in computing. We do this by responding to four key questions raised by Dr. Dennis Bushnell at the conclusion of our Creativity and Innovation presentation on August 1st entitled: "*Computing Faster Without CPUs*".

---

1. How does ultrafast computing research help Langley & NASA missions?

(Is it important to achieve computing speeds at and beyond Teraflops for engineering applications?)

A review of the NASA, Langley and branch mission and vision statements identify the following highlighted (bold) goals the products of our research promise to advance:

NASA's Vision for the Future
To **improve life here**
To **extend life to there**
To **find life beyond**

Langley Mission:

In **alliance with industry**, **other agencies, academia** and the atmospheric research community, in the areas of aerospace vehicles, aerospace systems analysis and atmospheric science we undertake **innovative, high-payoff activities beyond the risk limit or capability of commercial enterprises** and **deliver validated technology, scientific knowledge and understanding** of the Earth's atmosphere.

Analytical and Computational Methods Branch Mission:

Perform research and technology development of **efficient, physics-based analytical and computational methods** to enable multidisciplinary **design and analysis of advanced materials and structures for aerospace applications**.

Crew Systems and Operations Branch Mission:

Explore technology opportunities that **enhance the growth of the global aerospace system.** This includes the **identification, definition, development, and evaluation of system concepts**, the development of **methodologies and guidelines toward the development of these products** and the development of **enabling operational solutions** that support these technology opportunities.

What these Missions mean to us

Looking from the BIG Picture down to our research, we see:

Complex NASA missions require advanced, "smart" vehicles and systems

v

Design of these missions, their execution and design of complex vehicles to fly these missions are enabled by and benefit from ultrafast computing

v

Ultrafast computing is enabled by synergy of the following developments:

New H/W + S/W (*a.k.a. Gateware*), reconfigurable computing architectures, and new solution algorithms in engineering applications. Our specific instantiation of the above is FPGA + VIVA + new computational methods for structural mechanics.

Consider word-processing as an analogy. When one writes a letter using a word processor on a PC, it is the writer's ability to express thoughts by sentences that paces the work. The hardware & software infrastructure has more than enough speed to keep up with even the fastest thinker.

Not so when the task is engineering design. Then, the "what if" questions may arise rapidly, but the answers come slowly. It may be hours, or even days, before answers to "what if's", frequent in a complex aerospace vehicle design, can be analyzed and answered. That slows the human creative train of thought to the detriment of the design quality and timeliness. The computer hardware & software speed paces the designer's work.

A Langley Committee postulated in "Compute as Fast as Engineers Can Think"[1] capability to enable design processes paced solely by the human ability to conceive ideas and ask the right "what if" questions.

The Langley Committee examined a sample of aerospace vehicle design scenarios and concluded that a sustained computing speed on the order of *Exaflops* (3 orders of magnitude past Teraflops) is needed to achieve the desired capability. FPGA's coupled with new methods to exploit them, offer potential for these ultrafast engineering computations.

NASA needs to be involved in reconfigurable computation research because it represents no less than the future of all computing. FPGAs are simply the current enabling technology. The real research questions lie in how do you design a computing device (whether it is a supercomputer, an avionics component, a sensor system, satellite, etc.) that has dynamic architecture resources. With reconfigurable computing, the line between hardware and software is blurred or non-existent. Much research has been done and remains to be done on how to control a dynamic resource for efficient operation. Reconfigurable Computers will likely require a new design paradigm that is different than current hardware and software design paradigms. The fact that hardware resources can be reconfigured during operation introduces a new degree of freedom that is not trivial to manage. But given the progress that has been made in this area with FPGAs it is clear that architectures are evolving in this direction. In the not-to-distant future it is not hard to predict that ALL computing devices will have reconfigurable elements. It is a natural progression of the technology that mimics nature with respect to living organisms.

2. Why NASA? Won't others do this research (100TFLOPs at DOE)?

(What further gains in computing speed may be expected considering that others reported speeds of the order of Teraflops?)

These speeds were data *processing* **peak** speeds, not the **sustained speeds** achievable when *solving* science and engineering problems. There is a fundamental difference between processing a large volume of data that are independent in the sense of not being coupled by equations, and an equally large set of data defining equations to be solved.

Processing large volumes of independent data occurs typically in telemetry, and visualization (digital movies) where the data processing is merely a type of data sorting. It is relatively easy, then, to achieve spectacular speeds by employing large number of CPU's operating on a data stream appropriately partitioned.

Not so when equation solving is the problem. Then, partitioning of the equations into simultaneously processed subsets, a subset per processor, exacts a time penalty for cross-processor communication so that the law of diminishing returns in regard to the aggregate speed sets in quite early. Experience with applications in structures and fluids shows that the speed vs. number of processors levels off at a value that is problem-dependent and typically equal to only a few hundred processors operating simultaneously. In this sense there is no scalability.

The above scalability barrier may be pushed back to some extent by development of new methods tailored to a simultaneous use of many processors but a new barrier is soon encountered, posed by the dynamic variability of the number of processors that can be gainfully employed concurrently. In other words, it is typical that the number of such processors that can operate concurrently varies as the solution advances through its algorithmic stages.

Here is where FPGA's constitute a true breakthrough: the number of processors one creates in an FPGA may not only be very large and expandable but it can also vary on the fly from one phase of the solution algorithm to the next as needed. That variance may, of course, change from one problem to another. In this sense, the FPGA parallelism is inherent and truly scalable.

That's why FPGA technology, when combined with new development methods tailored to exploit its reconfigurability, has the potential of moving science and engineering computing past TeraFLOP speeds, and achieving a sustained, as opposed to the peak, performance.

Recognizing a high-performance computing gap (below) DARPA recently initiated a decade long (3 phase) program aimed at national security (cryptography) and military needs. Unfortunately, the DARPA phase 1 selectees exclude reconfiguable computing and FPGA-based innovations, leaving NASA in the leadership role:

"The High Productivity Computing Systems (HPCS) program is pursuing the research and development of viable high productivity computing system solutions that will fill a DoD high-end computing ***gap between today's late 80's based technology High Performance Computing (HPCs) and the promise of quantum computing****. DARPA's 'Grand Challenge" is to develop a broad spectrum of innovative technologies, integrated into a balanced total system solution by the end of this decade. The end product will be economically viable high productivity computing systems with both scalable vector and commodity HPC system functionality for the national security and industrial user communities…"* (See: http://www.darpa.mil/ipto/Solicitations/PIP_02-09.html)

Our Langley research plays a key leadership role at the forefront of reconfigurable computing and FPGA innovation focused on scientific and engineering applications.

3. Can Langley make an impact when industry spends $100B/yr on IT research?

(What difference will the NASA funding of a few hundred thousand dollars make, versus the billions of R&D dollars the hardware and software industry is spending already?)

The computer hardware & software industry is, indeed, spending billions but hardly any of that funding relates to the development postulated by the preceding two answers. Nearly all that funding goes into hardware and software aimed at the broad market of public and business utility. That market size is so much larger than the market of scientific and engineering computing that provides the seed money. Paraphrasing O'Keefe's recent pronouncements, NASA should see it as apart of its mission to act where private sector sees the risk to high or returns too far off. New methods development coupled with new, unconventional hardware fits that point exactly and it focuses on a niche currently neglected and unlikely to attract another sponsor and is likely to lay fallow unless NASA takes the initiative.

According to the NRC study "Evolving the High Performance Computing and Communications Initiative to Support the Nation's Information Infrastructure:

http://www.nap.edu/readingroom/books/hpcc

"The government-supported research program (on the order of $1 billion for information technology R&D) is small compared to industrial R&D (on the order of $20 billion), but it constitutes a **significant portion of the research component**, and it is a **critical factor** because it supports the **exploratory work that is difficult for industry to afford**, allows the pursuit of ideas that may lead to success in unexpected ways, and nourishes the industry of the future, creating jobs and benefits for ourselves and our children. The industrial R&D investment, though larger in dollars, is different in nature: it focuses on the near term -- increasingly so, as noted earlier -- and is thus vulnerable to major opportunity costs." (Page 24)

Langley C & I (Year 1) Hypercomputer research[2-12] results fit this goal by exploring a new, innovative area that industry, with its focus on near-term results, has largely overlooked.

4. How does this relate to nano- and quantum computing?

(Why should we invest in FPGA-tailored methods considering that the nano and quantum-computing technologies may supercede FPGAs?)

Regarding the nano-technology, its impact on computers will be in further miniaturization with the attendant speed gains to accrue from shorter data pathways. That impact will affect FPGA just as any other hardware type. There is nothing fundamentally different when computers go nano.

In contrast, quantum computers are very different. They are inherently parallel in a two-fold sense as they enable concurrent operation with an astronomical number of concurrent operations and, beyond that, their hardware elements may exist simultaneously in many different states.

One may consider FPGAs as a precursor to quantum computing with regard to the large number of concurrent operations but not in regard to simultaneity of states – the latter being a uniquely quantum feature. That means that the methods developed for and tried on the FPGAs that exploit large numbers of concurrent operations should be transferable to the quantum computers when these computers mature into application-ready tools in due course (20 years?), while producing benefits to the users relatively soon.

Quantum Computing Outlook of experts Drs. LaFlame and Luo

On June 27th Drs. Storaasli and Sobieski conducted a phone interview (45 minutes) with Dr. Ray LaFlame, developer of perhaps the most advanced Quantum Computer to date (the 7-qbit quantum computer at Los Alamos National Labs). Ray completed his doctorate at Cambridge University as a student of Stephen Hawking. Despite the remarkable progress he achieved at Los Alamos (funded by NSA), Ray indicated Quantum Computers now and in the future would be applicable for a small subset of computing applications, namely for cryptographic and quantum simulation. As far as when such a quantum computer might be developed to solve even this subset of applications, Ray said definitely NOT in 10 years and it would even be questionable for 20 years. As regards NASA type applications in Structures, Fluids or Electromagnetics or general Classiclal Mechanics and matrix solution **[A]{x} = {b},** Ray said it didn't look possible to him although it's possible with a major breakthrough something might be achievable in 20 years. On the positive side he said of his very basic research: "Right now it's impossible to say if we can scale these technologies, but if you asked me 5 years ago if we could build a 7-qubit computer in 5 years, I would have said it was impossible". In followup discussions with Dr Luo (ICASE) who worked down the hall from Dr LaFlame at Los Alamos before joining ICASE, Dr. Luo confirmed all of what Ray LaFlame related to us. Further, he added that he thought our research in harnessing reconfigurable computers (whether FPGA or followons) appeared to him to show more promise in the next 10-20 years for ultrafast computing for science and engineering. He has requested to join us in our ultrafast computing research. Regarding the leaders in the world in Quantum computing besides himself, Ray included: ATT (Shor), Caltech (Kimble), UC, LANL (Hughes), MIT (Lloyd, Cory), NIST (Wineland), Harvard (Havel), Stanford (Harris), the UK and Switzerland. Dr Luo noted that a Cornell Professor has an excellent web site to get up to speed on quantum computing:

http://www.ccmr.cornell.edu/~mermin/qcomp/CS483.html

Drs. Storaasli and Sobieski shared this and their Quantum Computing research proposal on July 2nd with Langley's Quantum Technology team led by Norm Barnes (who reported it in the minutes).

According to Kent Gilson, Chief Technologist at Star Bridge Systems "*Every new computing technology, including Nano-computing and Quantum Computing, is likely to find its initial incarnation as FPGAs. Viva will retarget to both of these substrates and, in fact, may be the only language that will*" (Aug2, 2002 email communication).

According to Kent Gilson on August 2nd, 2002: " It will possibly be **10 years until Nano** Computers may be serious HPC contenders in HPC market and at least **20 years until Quantum** Computers may be serious contenders in HPC market. There are fundamental problems involving the laws of physics to be faced that are not easily overcome (and significant funding to occur) before Quantum Computers can be a reality".

NSA (Hunsberger) Quantum Computer Outlook

On August 5th, Olaf Storaasli received the following E-mail from Alan Hunsberger who leads an innovative High-Speed Computing Group at the National Security Agency (Ft. Meade). Alan and his group have visited with members of our Langley group on at least three occasions prior to NSA purchase of a number of the latest Hypercomputers from Star Bridge Systems for Cryptography. Through NSA's LUCITE program, three additional Star Bridge Hypercomputers were recently installed at GWU, GMU and USC. NSA appreciates NASA leadership in this new technology as NASA has acted as a "testbed, proof-of-concept and demonstration" with NASA "spinoff" to NSA.

From: "Alan, Hunsberger" <awhunsb@afterlife.ncsc.mil>
Date: Mon Aug 05, 2002 08:20:30 AM US/Eastern
To: Olaf Storaasli <olaf@cox.net>
Cc: tom page <page@nsa.gov>, "Hunsberger, Alan," <awhunsb@afterlife.ncsc.mil>
Subject: Re: Quantum and Nano computer on the "Horizon"
Reply-To: awhunsb@afterlife.ncsc.mil

Olaf,

A group from government, academia and industry recently completed a submittal to Congress outlining a long-range high-end computing program. The document is still being reviewed and I don't think it has been submitted to Congress yet. Tom Page (cc above) led the technology subgroup and may have some input for you also. He may be able to send you a copy of the report. There was a member of the operation users group from NASA Ames, but I don't know the name of the person.

The basic conclusion on QC is: **"Quantum computers will not displace conventional high performance computers.** There are only a **very few problems for which theoretical QC algorithms are known to produce exponential speedup.** If feasible, and barring unforeseen breakthroughs, QC of the scale necessary to attack even those limited problems at a level **useful for national security** is at least **20 years away.** QC should therefore be considered outside the scope of this program." It is envisioned that QC will be used as a co-processor: "...the most likely outcome [for QC] is the creation of a third type of HPC co-processor system [clustered symmetric multiprocessor and clusters of scalable vector processors being the other two] that may be the system of choice for a small set of problems." "The paradigm of an HPC system as an adjunct to a QC or (more likely) a QC as an adjunct to an HPC system is the most likely QC paradigm. In any event, it is unlikely that quantum computers will replace any classical HPC systems."

Nanoelectronics for logic, memory and on-chip interconnects was viewed as a potential for deployment in **10+ years.** (given proper support for R&D).

I hope the above helps. Have a great vacation.

Alan

References

1. **Biedron, R.T.; Mehrotra, P.; Nelson, M.L.; Preston, F.S.; Rehder,J.J.; Rogers Jr., J.L.;Rudy, D.H.; Sobieszczanski-Sobieski, J.; and Storaasli, O.O.: Compute as Fast as Engineers can Think!;** *NASA/TM-1999-209715.*

2. **Singleterry, Robert C., Jaroslav Sobieszczanski-Sobieski, and Samuel Brown. "Field-Programmable Gate Array Computer in Structural Analysis: an Initial Exploration."** *43^{rd} American Institute of Aeronautics and Astronautics (AIAA) Structures, Structural Dynamics, and Materials Conference*. April 22-25, 2002.

3. **Storaasli, Olaf O., Robert C. Singleterry, and Samuel Brown. "Scientific Computations on a NASA Reconfigurable Hypercomputer."** *Abstract accepted for 5^{th} Military and Aerospace Programmable Logic Devices (MAPLD) Conference, Paper in preparation*. September 10-12, 2002.

4. **Fithian, William, Samuel Brown, and Tyler Reed. "Object Synchronization in VIVA 1.5."** *Briefing prepared for VIVA users at NASA Marshall, Eglin AFB, Progress Forge, Inc., and Star Bridge Systems, Inc.* March 26, 2002.

5. **Barr, Kristen, Shaun Foley, and Robert A. Lewis II**. "Hypercomputing with the CORDIC Algorithm." August, 2001. Presentation of research conducted under Dr. Olaf O. Storaasli, June-August, 2001.

6. **Butler, Patrick**. New Horizons Governors School Mentorship Project. May, 2001. Presentation of research conducted under Dr. Olaf O. Storaasli, September 2000 – May 2001.

7. **Dandawate, Neha.** "Reckless Speeding: The Investigation of the Programming Capabilities of the HAL Hypercomputer." July, 2002. Presentation of research conducted under Dr. Olaf O. Storaasli, June – July, 2002.

8. **Dandawate, Neha.** "The Investigation of the Programming Capabilities of the HAL-15 Hypercomputer." July, 2002. Paper on research conducted under Dr. Olaf O. Storaasli, June – July, 2002.

9. **Fithian, William.** "Developing a Matrix Equation Solver for the HAL-15 Hypercomputer." December, 2001. Proposal for research to be conducted under Dr. Olaf O. Storaasli, September 2001 – May 2002.

10. **Fithian, William.** "Developing a Matrix Equation Solver for the HAL-15." May, 2002. Presentation of research conducted under Dr. Olaf O. Storaasli, September 2001 – May 2002.

11. **Foley, Shaun.** "Scientific Hypercomputing." August, 2001. Paper describing research conducted under Dr. Olaf O. Storaasli, June – August, 2002.

12. **Krishnamurthy, Siddhartha.** "Development of an Integration Algorithm for Field Programable Gate Arrays using VIVA." July, 2002. Paper describing research conducted under Dr. Robert C. Singleterry, June – Aug 2002.

Reconfigurable Computing Research Awards:

Patrick Butler: **Best Mentorship Project** (awarded by New Horizons Governor's School for 2000-2001 research)

William Fithian: **Oracle Award** for Computer Science (one of nine national scholarships), **Best Mentorship Project** (awarded by New Horizons Governor's School for 2001-2002 research).

Further Information: **http://hummer.larc.nasa.gov/acmbexternal/Personnel/Storaasli/reconfig.html**